

Please replace the paragraph beginning on page 7, line 4 with the following rewritten paragraph:

C3
-- Referring now to FIG. 2, the equipment interface layer 16 of the accounting process 14 includes various equipment interfaces 42a-42c which are, respectively, an interface 42a for the router/switch 12a, an interface 42b for the cable/modem head end 12b, and an interface 42c for the flow probe 12c. The equipment interface layer 16 also includes additional interfaces such as an interface 42d for a remote access concentrator 12d, an interface 42e for an Extranet switch 12e, an interface 42f for a DNS server 12f, and an interface 42g for a RADIUS server 12g. The equipment interface can have additional interfaces that can be specified, as new equipment is added. The interfaces 42a-42g can be developed by an interface toolkit 44. The interface toolkit 44 permits a user to construct a new equipment interface type to couple the accounting process 14 to a new equipment source type.--

Please replace the paragraph beginning on page 8, line 24 with the following rewritten paragraph:

C4
-- The accounting process 14 also includes a flow aggregation processor 60 that is part of the aggregation and distribution process 17 (mentioned above). The flow aggregation processor 60 is a central collection point for all network accounting records (NARs) produced from various data collectors 52a-52g in the flow data collection layer 18. The flow aggregation processor 60 receives NARs from various data collectors 52a-52g and aggregates, i.e., summarizes related information from the received NARs across the accounting support arrangement 10. The aggregation processor 60 produces Summary NARs i.e., enhanced and unique network accounting records. That is, the flow aggregation process aggregates the records across the network devices; whereas, individual data collectors 52a-52g can aggregate accounting records from individual data sources. Aggregation will be described below in FIGS. 14-23.--

Please replace the paragraph on beginning on page 10, line 2 with the following rewritten paragraph:

C5
-- As shown in FIG. 3, for the Internet service provider, data collectors 52a-52d (illustrated in Fig. 2) are distributed at specific Points of Presence (POP), such as remote access concentrators 102 managed by the Internet service provider. The remote access concentrators

C5 allow a mobile Internet user 106 or an Internet user 107 with remote access to access an enterprise over the Internet, via the Internet service provider. In this example, the Internet service provider arrangement 100 and the large Enterprise arrangements 110 and 120 include servers 13, 13', and 13'' that run accounting processes 14, 14' and 14''. The accounting processes 14, 14' and 14'' each independently manage and collect information regarding network traffic usage--.

Please replace the paragraph beginning on page 12, line 8 with the following rewritten paragraph:

C6 -- Referring now to FIG. 4, a similar access configuration 100', as the configuration 100 (FIG. 3) can be used with an Extranet switch 122. Extranet access allows remote users to dial into an Internet service provider (ISP) and reach a corporate or branch office via an ISP. The Extranet switch 122 allows Internet users access to corporate databases, mail servers and file servers, for example. It is an extension of the Internet in combination with a corporate Intranet. In this configuration, the Extranet switch 122 can be owned and operated by an Internet service provider as shown with enterprise A, or it could, alternatively, be owned and operated by an enterprise, as shown with enterprise B. Users would access the corporate network of either enterprise A or enterprise B, via the Internet service provider with various types of tunneling protocols such as Layer 2 Tunneling Protocol (L2TP), Layer 2 Forwarding (L2F), Point to Point Tunneling Protocol (PPTP) or Internet Protocol Security (IPSec), and so forth. The accounting server 13 located at the service provider and also accounting servers 13', 13'' within enterprise A and enterprise B allow each the Internet service provider and each of enterprises A and B to run accounting process 14', 14'' on the servers 13', 13'' to monitor and collect network data.--.

Please replace the paragraph beginning on page 13, line 1 with the following rewritten paragraph:

C7 -- Referring now to FIG. 5, a graph 140 depiction of a very large scale network includes a device "A" 142 communicating with a device "B" 144. The graph 140 includes nodes (not all numbered) that can represent routers, switches, flow probes, etc. that have interfaces (not shown) which maintain statistics on information passed through the interfaces. For example, a switch may have a number of Ethernet ports and a host could be connected to one of the ports and in

C7 communication with one of the interfaces to transfer information over the network. The interface would have counters that are used to track "packets in," "packets out," "bytes in," "bytes out," and so forth.--

Please replace the paragraph beginning on page 13, line 12 with the following rewritten paragraph:

C8 -- In this case of the host connected to the port, or a router or some other device being connected to the port, there is no other connection that the host, router or other device is aware of other than the entire network. This is an example of a "connectionless oriented" protocol. A data collector 52 can be disposed in the network in a path between the entities "A" and "B", such that the data collector 52 monitors some of the packets that comprise a flow between "A" and "B." As a single point monitor, the data collector 52 has no concept that there are two ends communicating. The data collector 52 identifies these entities "A" and "B" in various NARs produced by the data collector 52. At a later stage in the processing, either in the data collector 52 or elsewhere in the accounting process 14 the NARs are correlated so that the NARs or some aggregated NAR produced by the data collector 52 or the rest of the accounting process 14 can be associated with the accountable entities "A" and "B" to thus identify a connection between entities "A" and "B."--

Please replace the paragraph beginning on page 14, line 23 with the following rewritten paragraph:

C9 -- Thus, the data collector 52 is a single point monitor that monitors traffic at one point in the network and converts the traffic into "pipe oriented" or "flow oriented" accounting information. The data collector 52 identifies a source and a destination of the traffic. That is, the data collector 52 develops a "connection oriented tracking." By distributing data collectors 52a-52g (FIG. 2) throughout the network, the network can be modeled as pipes having two endpoints. A data collector can be disposed in a partial pipe. The data collector 52 determines that one end of the pipe refers to "A" and the other end of the pipe refers to "B." The data collector 52 can be disposed anywhere along the network.--

Please replace the paragraph beginning on page 15, line 13 with the following rewritten paragraph:

C10
-- Some equipment have a half pipe model that generates independent accounting records for each half pipe. The data collectors can assemble full pipe information from half pipe information. The accounting process 14 could be coupled to equipment that gives a half pipe model for A communicating with B and a separate one for B communicating with A. The data collectors 52a-52g combine information from these two half pipes into a bidirectional flow.--

Please replace the paragraph beginning on page 15, line 21 with the following rewritten paragraph:

C11
-- Referring now to FIG. 6, an example of data flow 130 through the accounting process 14 is shown. In this example, the data flow 130 is initiated by a user 131 making a call to a remote access concentrator (RAC) 132. Upon receiving the call, the RAC 132 authenticates the user 131 against a secure access controller 134. After verification, the RAC 132 connects the user 131 to the network 135 and sends a RADIUS Start record (not shown) to the accounting process 14. The accounting process 14 generates a RADIUS Start NAR 137a and stores the RADIUS Start NAR 137a in a database 62. At that point, the remote user may check e-mail, look at a web server and transfer a file. For each transaction, the accounting process 14 captures the IP traffic, generating e-mail, http, and ftp network accounting records 137b-137d, respectively. These are stored in the database 62. Upon completion of these transactions the user would log out of the network, at which time the RAC 132 would send the accounting process 14 a RADIUS Stop record. The accounting process 14 generates a RADIUS Stop NAR 137e and stores the RADIUS Stop NAR 137e in the database 62. All of these records reflecting the user's transactions could be viewed and reported in flexible ways dependent on the needs of an end-user application.--

Please replace the paragraph beginning on page 16, line 27 with the following rewritten paragraph:

C12
-- FIG. 7 has a plurality of exclusively "Activity NARs" which could correspond to a very low level of detail, or could be the result of a prior aggregation providing a higher level view of the information. Thus, FIG. 7 shows a collection 152 of exclusively activity NARs. From base

C12
level data, additional "views" of the NAR could be produced, such as a set of "Summary NARs" 154, or another set of Activity NARs 156 which could be a result of further aggregation of the base level information, or lastly a combination of a set of Summary NARs and Activity NARs 158. The summary NAR is produced by the central aggregation processor 60 and can include user identifying information, protocol information, connection time information, data information, and so forth.--.

Please replace the paragraph beginning on page 21, line 13 with the following rewritten paragraph:

C13
-- The plurality of Network Accounting Record Attributes 204a-204n provide metrics for the NAR 200. The Network Accounting Record Attributes 204a-204n capture specific information contained in data from network devices. Differentiating between the Network Accounting Record Identifier 202 and the metric 204 allows the accounting process 14 to perform logical and arithmetical operations on metrics 204 while leaving the Network Accounting Record Identifier 202 intact. The Network Accounting Record Identifier 202 can be enhanced unlike the metrics 204.--.

Please replace the paragraph beginning on page 21, line 22 with the following rewritten paragraph:

C14
-- The data collectors 52a-52g (FIG. 2) are oriented around the process of filling in the NAR. The metrics are left untouched by the data collector and are passed transparently into the accounting process flow aggregation processor 60. The data collectors 52a-52g assign the Network Accounting Record Identifiers 202 to the metrics e.g., a source and a destination identifier to the metric. In the example of a router link, the metrics that the router interface provides are in the form of "information in" and "information out" e.g., octets in, octets out, bytes in, bytes out, datagrams in, datagrams out, faults in, faults out, and so forth. The data collectors 52a-52g determine what "in" and "out" mean and assigns the unique identifier that is unambiguous relative to the determined meaning of "in" and "out." Once a data collector 52 has established this convention, the convention is used throughout the system 10.--.

BEST AVAILABLE COPY

Please replace the paragraph beginning on page 33, line 17 with the following rewritten paragraph:

C15
-- Referring now to FIG. 15, a data collection process 330 performed by the flow data collector 52 of FIG. 17 is shown. The flow data collector receives 332 data from the equipment interface for a network device. The flow data collector performs an equipment interface specific translation to convert 336 the received data into NAR format as well as populates the NAR header. Once the NAR is populated with the appropriate data, the flow data collector 52 attempts to correlate 338 the newly populated NAR with the other NARs. That is, the flow data collector 52 compares the newly populated NAR to NARs currently stored in the local store 314 (from FIG. 14) to determine if there are multiple instances of the same object. Specifically, correlation is performed by examining the ACCT_ENTITY_ID (from FIGS. 11A-11E).--.

Please replace the paragraph beginning on page 37, line 27 with the following rewritten paragraph:

C16
-- The flow aggregation processor (FAP) 60 (FIG. 2) aggregates and/or enhances record data across the system 10. It receives data from multiple flow data collectors (FDCs) that may be aggregating and enhancing close to the source of the information (as described above with reference to FIG. 17). As NARs are received from multiple FDCs, the data can be further enhanced and/or reduced (i.e. aggregated) to meet the specific needs of an application or output interface based on the aggregation policy of the flow aggregation processor 60 (FAP). The design and operation of the FAP will be described in more detail below.--.

Please replace the paragraph beginning on page 41, line 6 with the following rewritten paragraph:

C17
-- These two records NAR1, NAR2 are combined through correlation 442 (from FIG. 17) and enhancement 444 (FIG. 17) to generate an enhanced NAR2 532. This enhanced NAR has a modified accountable entity identifier 534 and a metric. The modified accountable entity identifier is the existing accounting entity ID 514, to which the FAP has added the IP-to-username assignment from the accounting entity ID 512 of the NAR1 508.--.

Please replace the paragraph beginning on page 41, line 13 with the following rewritten paragraph:

C18 -- Still referring to FIG. 18, the NAR1 508 has an IP-to-username mapping 512 and an accounting interval 516 comprising a start time and a session time to indicate a time interval bounded by start time "T1" and a start time + session time ("T2"), that is, the accounting interval represents a start time and a stop time. The username 524 in the IP address-to-username mapping is supplied by the DHCP server 500. In the FAP, this NAR1 information will either go directly to a correlation function or to the local store (which could either be a database, file or memory), where it can be directly accessed by the correlator function. The NAR2 510 has an accounting entity ID 514, a T3-to-T4 accounting time interval 518 and a metric 530. The accounting entity identifier 514 has two IP addresses 526, 528, one corresponding to a source IP address and the other corresponding to a destination IP address. The NAR2 510 is passed to the correlator 442, which determines that the T1-to-T2 time interval 516 from the IP-to-username address map in the NAR1 508 overlaps or in some way relates to the T3-to-T4 time interval 518 of the NAR2 510. The correlator 442 determines that T1, T2, T3 and T4 are related, and that the IP address 522 in the IP-to-username address mapping 512 is associated with one of the two IP addresses 526, 528 in the NAR2 510. Thus, the FAP enhances the NAR2 510 by inserting information from the accounting entity ID 512 (of NAR1 508) into the accounting entity ID portion of the NAR2 510. The resulting, enhanced NAR2 532 has an enhanced accounting entity ID 534 that includes the T3-to-T4 timestamp (not shown), the IP-to-IP addresses 526, 528 and the username 524. Thus, the enhanced NAR2 now has a mapping between the username and the one of the IP addresses 526, 528 that is related to the IP address 522. The metric 530 is unchanged.--

Please replace the paragraph beginning on page 49, line 20 with the following rewritten paragraph:

C19 -- As discussed above in reference to FIG. 2, the accounting process supports a flow probe e.g., 12c that captures a user's network activity for purposes of IP accounting. The flow probe 12c monitors all traffic over a given network link and captures data associated with the different flows in the traffic on that link. It is capable of monitoring IP data flows over a number of technologies (e.g., Ethernet, Asynchronous Transfer Mode (ATM), FDDI, etc.).--.

Please replace the paragraph beginning on page 51, line 9 with the following rewritten paragraph:

C20 -- Generally, a flow is defined as any communication between communicating entities identified by an IP address, a protocol and a service port. All IP packets (or datagrams) are categorized using the fields present in the packets themselves: source/destination IP addresses, the protocol indicated in the IP header PROTO field, and, in the case of User Datagram Protocol (UDP) or Transmission Control Protocol (TCP), by the packet's source and destination port numbers.--

Please replace the paragraph beginning on page 58, line 20 with the following rewritten paragraph:

C21 -- The flow probe reports on network traffic activity through a flow probe NAR, which reports IP flow traffic activity. The flow probe categorizes network traffic into one of four classes of traffic flow: i) connection oriented (e.g., TCP); ii) new connectionless; iii) request/response connectionless (e.g., User Datagram Protocol (UDP), Domain Name System (DNS)); and iv) connectionless persistent (e.g., Network File System (NFS), Multicast BackBONE, or "MBONE" multicast traffic). To each of these classes it applies connection oriented semantics for a uniform approach to status reporting. That is, the flow probe treats these dissimilar transaction models as if they were the same. There is one uniform structure for the status reports generated for each of the 4 different transactions. Each status report includes transaction start and stop information, media access control (MAC) and IP source and destination addresses, the IP options that were seen, the upper layer protocol used, and the transaction source and destination byte and packet counts and upper layer protocol specific information. The protocol specific information and the criteria for when the status reports are created is different for each of the four transaction types.--

Please replace the paragraph beginning on page 63, line 23 with the following rewritten paragraph:

C22 -- For some protocols that permit wrap around, the packet loss detector process 704 tests 718 if the sequence number has wrapped around e.g., gone from 32 bits of all ones to 32 bits of

C 22
all zeros. The IPSec Authentication packets currently do not permit wrap around, so test 718 would not be necessary for IPSec Authentication Headers. If for other protocols (or latter versions of the IPSec Authentication protocol), the packet loss detector process 704 detects a wrap around condition, then there has not been any packet loss and the packet is dropped. The packet loss detector process 704 will update 712 the stored sequence number for that flow in the cache. If the sequence number is any other number, i.e., it did not turn over to all zeros, then there may have been packet loss. If there may have been packet loss, the packet loss detector process 704 can determine how many packets have been lost by determining how many sequence numbers are missing.--.

Please replace the paragraph beginning on page 66, line 15 with the following rewritten paragraph:

C 23
-- An important component of quality of service includes determining whether there has been packet loss. The packet detector monitor described in conjunction with FIGS. 29A and 29B can be used to access packet loss. The packet detection monitor 702 can be deployed in the network and generate NARs that can be used to determine packet loss as discussed above. This information can be used in the capturing quality of service process 730 to assess whether the policy specified by the service level agreement was provided to the customer. Additionally, so called Differentiated Service "DiffServ technology" that a known quality of service solution that has been proposed for the Internet as well as enterprise networks. In contrast to a per-flow orientation of some types of quality of service solutions such as Integrated Services (Int-serv) and Resource Reservation Setup Protocol (RSVP), Diffserv enabled networks classify packets into one of a small number of aggregated flows or "classes", based on bits set in the type of service (TOS) field of each packet's IP header. This quality of service technology for IP networking is designed to lower the statistical probability of packet loss of specific flows. The capturing quality of service process 730 establishes DiffServ policy that is decomposed into a collection of DiffServ configurations. The DiffServ configurations are deployed to a collection of routers or switches that the customer would have access to in the network 11 as part of the enforcement/deployment process 732. Because packet loss is a statistical phenomenon, the capturing quality of service process 730 observes 736 a large number of network flows. The capturing quality of service process 730 can observe network traffic because of the use of the

C23 accounting process 14 and the resulting NARs at the granularity in which the DiffServ policies are actually being deployed. The DiffServ policies are generally deployed at the source and destination IP address, protocol and possibly destination port level.--

Please replace the paragraph beginning on page 68, line 19 with the following rewritten paragraph:

C24 -- A service management feedback process 750 therefore includes three components, service provisioning 752, policy server 754 and service accounting 756. The role of service provisioning 752 is to send requests 752b to the policy server 754 to obtain an appropriate active policy, and to obtain rules and domain information 754a from the policy server. The provisioning system can communicate with appropriate network management systems and element management systems (not shown) to configure the network 10 for an end-to-end service. When the configuration 752a is deployed at the various network devices (not shown), the service is produced. The level of service is monitored or audited by the accounting system 756 which can be the accounting process 14 described above. The accounting process 14 monitors the level of service by producing appropriate network accounting records. The network accounting records (NARs) are used by a billing application to adjust billing based on the level of service that was provided as determined by the accounting process 14. The accounting process 14 also can compare the policies produced by the policy server to the actual levels of service provided to the customer by examining NARs that are produced by the customer's usage of the network.--

Please replace the paragraph beginning on page 69, line 12 with the following rewritten paragraph:

C25 -- In addition, levels of service might change, and the system takes changes into account so that the service management can modify the charge or account differently for those changes in levels of service. The service accounting also uses the active policy information from the policy server to deliver billing information to a billing system or to a chargeback system that can make adjustments to billings for the service.--

Please replace the paragraph beginning on page 69, line 19 with the following rewritten paragraph:

BEST AVAILABLE COPY

C26 -- A policy server 754 is built on the capabilities of address management, domain name management and so forth. Essentially in a policy enabled network, policy servers produce a set of rules and apply those rules to a domain or problem set. The policy server communicates the rules to the accounting process 14 so that the accounting process 14 can determine what kind of records to generate. All of the information is described using data flows.--

Please replace the paragraph beginning on page 69, line 27 with the following rewritten paragraph:

C27 -- As an example, a service contract may specify that a company "X" will be given 100% availability of a particular network device e.g., a router (not shown) and its corresponding service. In order to assure that level of service, the policy server 754 sends that requirement in a template to the provisioning service 752 to produce a configuration file 752a to configure the router to give company "X" preferred use of the router. Therefore, every time a packet from company "X"'s network comes across the router, the packet will always be transmitted unless there is something wrong with the router. This may occur even if a packet of company "Y", which has a lower service level than company "X", is waiting in the router to be transmitted. The packet from company "Y" will wait because company "Y" is not paying for the quality of service that company "X" is paying for.--

Please replace the paragraph beginning on page 70, line 13 with the following rewritten paragraph:

C28 -- In that case, the provisioning service 752 configures the policy enforcement mechanism that was put into the router in the network. How the policy was defined to the provisioning equipment is that there is a one-to-one relationship between the policy and what the accounting process 14 will monitor in the network. The accounting process 14 will be aware that company "X" contracted to have 100% availability from the router.--

Please replace the paragraph beginning on page 70, line 20 with the following rewritten paragraph:

C29 -- The accounting process 14 will then take every source of information it has available and will construct an accounting record that reflects the level of service actually delivered to

BEST AVAILABLE COPY

C29

company "X." The accounting records produced are relative to two components, i.e., the router and the customer. The accounting process 14 is flexible and can generate accounting records of any flow abstraction. In the service management feedback process 750, the policy server 754 sends a flow based policy to the provisioning service 752. The provisioning service 752 uses a flow based policy to configure the network. That same flow based policy is passed to the accounting process 14, which can generate NARs having metrics that can be used to match the same level of those flows. The output of the accounting process 14 will determine whether the quality of service, availability, etc. that was contracted for in the contract 751 was provided. Therefore, the service management feedback process 750 provides the level of service that was delivered at the same semantic level as the actual contract.--

Please replace the paragraph beginning on page 71, line 9 with the following rewritten paragraph:

C30

-- Capturing quality of service as audited by the accounting process 14 includes detecting packet loss, as mentioned above. Each of the components managed by the service management feedback process 750 requires information. Therefore, the provisioning service 752 has to provision these various quality levels. The policy server 754, thus, keeps what is essentially enforcement of the levels of quality that are offered by different service types, and the accounting process 756 detects, monitors and audits whether those classes in quality of service are being delivered.--

Please replace the paragraph beginning on page 71, line 19 with the following rewritten paragraph:

C31

-- Referring to FIG. 32, an implementation of the provisioning service 752 is shown. The provisioning service 752 extends concepts of device management and network management into a service management layer of functionality. The provisioning service 752 includes a provisioning core 782, provisioning modules 784, and element managers 786. The provisioning service 752 is user focused rather than network focused, as in conventional network management. Network management involves communication with network systems and equipment. The provisioning service 752 is oriented more towards a user and a user's concepts of services. The provisioning service 752 provides an additional layer of abstraction that relates

92

C

C31 the description of services at a user level to a network's ability to provide those end-to-end services. The architecture 780 of provisioning service 752 is a multi-device 788 at the bottom of the architecture and multi-service 790 at the top of the architecture. The provisioning service 752 is deployed to write commands to the network systems, i.e., multi-devices 788 in order to change the configurations of those systems.--.

Please replace the paragraph beginning on page 72, line 9 with the following rewritten paragraph:

C32 -- Since many end customer services now require that a network operate with multiple kinds of network elements in order to provide an end-to-end service, the provisioning service 752 simplifies producing information that is necessary for a service provider to translate a service order from a customer to a network configuration, i.e., all commands necessary for all the different elements in the network in order to create an end-to-end service.--.

Please replace the paragraph beginning on page 72, line 17 with the following rewritten paragraph:

C33 -- The provisioning service 752 builds on existing systems. That is, in the lower layers there are existing element managers that have a configuration management system to configure at the network layer. The provisioning service 752 adds layering over the conventional network management layer. The provisioning service 752 maps a customer specified end-to-end service to the network elements that are required to produce that end-to-end service. Mapping of a customer's service orders into the state of the network can have various pieces of workflow necessary to create or completely activate this service order.--.

In the claims:

Please amend claims 1 as follows:

C34 1. (twice amended) A method of transmitting accounting records in an accounting system that produces information pertaining to network traffic flow comprising:

collecting data from a network device by a data collector associated with the network device and producing accounting records from the data;

BEST AVAILABLE COPY

C 34

transmitting the accounting records to first and second flow aggregation processes, with transmitting further comprising for each flow aggregation process:

- storing in the data collector the accounting records;
- transmitting the accounting records to the flow aggregation process; and
- awaiting an acknowledgment signal from the flow aggregation process that the flow aggregation process received the accounting records before discarding the accounting records sent to the flow aggregation process.

Attached is a marked-up version of the amendments made to the application by the current response. The attachment is captioned "VERSION WITH MARKINGS TO SHOW CHANGES MADE."

BEST AVAILABLE COPY